

Flux-Limited Diffusion for Multiple Scattering in Participating Media

D. Koerner¹, J. Portsmouth^{†2}, F. Sadlo¹, T. Ertl¹, and B. Eberhardt³

¹University of Stuttgart, Stuttgart, Germany ²The Moving Picture Company ³Stuttgart Media University

Abstract

For the rendering of multiple scattering effects in participating media, methods based on the diffusion approximation are an extremely efficient alternative to Monte Carlo path tracing. However, in sufficiently transparent regions, classical diffusion approximation suffers from non-physical radiative fluxes which leads to a poor match to correct light transport. In particular, this prevents the application of classical diffusion approximation to heterogeneous media, where opaque material is embedded within transparent regions. To address this limitation, we introduce flux-limited diffusion, a technique from the astrophysics domain. This method provides a better approximation to light transport than classical diffusion approximation, particularly when applied to heterogeneous media, and hence broadens the applicability of diffusion-based techniques. We provide an algorithm for flux-limited diffusion, which is validated using the transport theory for a point light source in an infinite homogeneous medium. We further demonstrate that our implementation of flux-limited diffusion produces more accurate renderings of multiple scattering in various heterogeneous datasets than classical diffusion approximation, by comparing both methods to ground truth renderings obtained via volumetric path tracing.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

1. Introduction

The rendering of multiply-scattered light is highly computationally expensive in the presence of dense participating media due to the proliferation of scattering paths. A brute force solution exists in the form of volumetric path tracing, in which the random scattering of photons is explicitly simulated via a Monte Carlo method such as introduced by Lafortune et al. [LW96]. In practice, however, these methods are computationally too demanding for many applications, since the calculation becomes increasingly expensive as the medium becomes more opaque and the mean free path (the local average distance travelled by a photon) between scattering events decreases. Eventually path tracing becomes intractable, and we enter a diffusive regime where photons undergo long random walks in the medium. In this regime, the classical diffusion equation provides a very good approximation to multiply-scattered light, which is vastly more efficient than path tracing. It is a simple partial differential equation which can be solved efficiently via a variety of methods, both analytical and numerical.

Diffusion methods are thus clearly of great interest as a means of accelerating the simulation of multiple scattering in high opacity media. However, application of these methods to more general volumetric media, such as clouds in which opaque regions are embedded in relatively transparent regions, has been problematic for the basic reason that the classical diffusion approximation (CDA) becomes increasingly inaccurate as the medium becomes more transparent. In fact, the classical diffusion equation for the light field becomes singular in the limit of a zero extinction “vacuum”

region, as elaborated on in Section 4. In the context of subsurface scattering [JMLH01], this is not an issue because the vacuum outside the medium is dealt with analytically by imposing boundary conditions at the surface. However, in a general heterogeneous volume this is not possible as there may not exist a well-defined boundary. So it has not been clear thus far how to extend the diffusion approach to these more general situations.

Our contribution is the introduction of the flux-limited diffusion (FLD) method to the graphics community and its application to the problem of rendering participating media with multiply-scattered light. This method is a standard radiative transfer technique in other fields (in particular, astrophysics [LP81, TS01] and nuclear physics [LK74]) and its underlying theory was first fully explored by Levermore and Pomraning [LP81]. FLD was developed specifically to produce a better approximation to multiple scattering in the low optical depth regime than CDA. We will explain how a grid-based diffusion solver can be easily extended to FLD by modifying the diffusion coefficient, and demonstrate that FLD in the mentioned cases performs significantly better than CDA, in the sense that it results in images closer to ground truth at affordable additional computational cost (see Figure 1). Beyond that, we believe that flux-limited diffusion has the potential to improve other techniques for multiple scattering in participating media which are based on diffusion approximation, such as [AWB11, ZM13].

In the next two sections, we discuss related work followed by a brief review of CDA theory. In Section 4, we introduce the theory of FLD and give some intuition for its physical basis. We show how to extend an existing CDA solver to FLD in Section 5. In Section 6 we evaluate our method using the analytical solution for a point source as a testbed, then render various datasets using FLD

[†] Joint first author




Figure 1: Renderings of a heterogeneous volume which contains regions of high and low opacity, embedded in vacuum. This is problematic for the classical diffusion approximation (c), which becomes invalid near and within regions of low opacity. Our approach employing flux-limited diffusion (b) addresses this deficiency and as a result provides a better approximation to volumetric path tracing (a) at a fraction of the cost.

and compare to ground truth images obtained via volumetric path tracing.

2. Related Work

Multiple scattering in participating media has been extensively studied and we restrict the discussion to the work which is most related to ours. For a broader overview, we refer to Cerezo et al. [CPCP*05].

The theory of multiple scattering as a diffusion process has a long history in the radiative transfer literature [Ish78, RL86]. In computer graphics, the theory of light diffusion has been most profitably applied to the rendering of subsurface scattering (SSS). This is the searchlight problem of finding the exitant radiance under incident illumination at the boundary of a homogeneous scattering medium. Jensen et al. [JMLH01] pioneered the application of the radiative diffusion approximation to this problem, solving the diffusion equation with appropriate boundary conditions imposed at the surface. Further work has refined this technique [DI11, YZXW12, HCJ13]. This line of research cannot be applied to our problem since SSS deals with media with well-defined surfaces, while we are interested in finding the radiance *within* heterogeneous volumes which contain regions of low opacity or vacuum.

Wang et al. [WZT*08] extended SSS to a heterogeneous translucent volumetric medium modeled as an irregular finite element grid on which CDA was solved. Arbree et al. [AWB11] later in similar approaches addressed various problems in Wang et al.'s work and simplified discretization by using a tetrahedral mesh basis. Li et al. [LSR*13] used the same tetrahedral representation but a more efficient method for constructing the coefficient matrix, which allowed realtime application under changing mesh topology. In contrast to our approach, all this work still focuses on rendering highly opaque media with well defined surfaces.

CDA was first applied to the multiple scattering of light in heterogeneous participating media defined on a grid by Stam [Sta95], refining earlier work by [KVH84]. Here, however, the breakdown of CDA within more transparent regions was not accommodated for. In applying CDA to the rendering of clouds, Max

et al. [MSM*04] recognized the problem, and to avoid it, introduced a hybrid scheme in which diffusion is applied to discrete homogenous cloud cores with high optical depth and well-defined boundaries on a grid. Ray tracing is used to propagate light across the empty space between grids.

Other related grid-based approaches to scattering in inhomogeneous participating media include Light Propagation Volumes [KD10, BCK*11], and Lattice Boltzmann methods [GRWS04, QXF*07]. These are similar to the diffusion methods in the sense that light is propagated between cells on a grid. Bouthors et al. [BNM*08] described a semi-analytical approach to multiple scattering in media bounded by geometry in which a pre-computed path tracing solution in slab geometry is used. All these methods are less physically-based than ours, using phenomenological models rather than direct approximation of the radiative transfer equation. Another general approach based directly on the radiative transfer equation (RTE) [Cha60] is Light Propagation Maps [Fat09]. This is essentially a formulation of the classic Discrete Ordinate Method [Cha60] in which the light field in each grid cell is discretized into ray direction bins. This is a potentially more accurate method than the diffusion approximation, but also much more computationally expensive.

3. Classical Diffusion Approximation

In this section we review CDA as a preliminary to our introduction of FLD in Section 4. For more details, we refer the reader to [Ish78] and [Sta95]. In Table 1 we give a summary of the symbols used.

We assume that the properties of the medium are specified, in the form of an absorption coefficient field $\sigma_a(\mathbf{x})$ and scattering coefficient field $\sigma_s(\mathbf{x})$, both with units of inverse length. We can also express these in terms of the total extinction field $\sigma_t(\mathbf{x}) = \sigma_a(\mathbf{x}) + \sigma_s(\mathbf{x})$ and the dimensionless scattering albedo $\alpha(\mathbf{x}) = \sigma_s(\mathbf{x})/\sigma_t(\mathbf{x})$. The mean free path is given by $\sigma_t(\mathbf{x})^{-1}$. We also allow for the possibility that the medium generates an isotropic radiance field $j(\mathbf{x})$ (with units of radiated power per unit volume) due to self-emission. We specialize in our whole treat-

Table 1: Nomenclature and units

Symbol	Meaning
σ_s	scattering coefficient [m^{-1}]
σ_a	absorption coefficient [m^{-1}]
$\sigma_t = \sigma_a + \sigma_s$	extinction coefficient [m^{-1}]
$\alpha = \sigma_s / \sigma_t$	albedo
$L(\mathbf{x}, \omega)$	radiance field [$\text{Wm}^{-2}\text{sr}^{-1}$]
$j(\mathbf{x})$	self-emission field [Wm^{-3}]
$Q(\mathbf{x})$	source term [$\text{Wm}^{-3}\text{sr}^{-1}$]
$L_{ri}(\mathbf{x})$	reduced incident radiance [Wm^{-2}]
$L_m(\mathbf{x})$	medium radiance [Wm^{-2}]
$\phi(\mathbf{x})$	diffuse fluence [Wm^{-2}]
$\mathbf{E}(\mathbf{x})$	diffuse flux [Wm^{-2}]
$D(\mathbf{x}) = \frac{1}{3\sigma_t}$	diffusion coefficient (CDA) [m]
$D_F(\mathcal{R}, \mathbf{x})$	diffusion coefficient (FLD) [m]
$\mathcal{R}(\mathbf{x})$	Knudsen number
$T(\mathbf{x})$	transmittance/DSM
τ	optical depth (number of mean free paths)

ment in this paper to the case of isotropic scattering and discuss anisotropy in Section 7.

Given these medium properties, the radiance field $L(\mathbf{x}, \omega)$, for all points \mathbf{x} in the volume and all ray directions ω , is in principle determined by solving the RTE, an integral equation which the radiance field must satisfy

$$(\omega \cdot \nabla)L(\mathbf{x}, \omega) = -\sigma_t(\mathbf{x})L(\mathbf{x}, \omega) + Q(\mathbf{x}), \quad (1)$$

with

$$Q(\mathbf{x}) = \frac{j(\mathbf{x})}{4\pi} + \sigma_s \int_{\Omega} \frac{1}{4\pi} L(\mathbf{x}, \omega') d\omega'. \quad (2)$$

The source term $Q(\mathbf{x})$ represents self-emission and in-scattering (the factor 4π in the first term is the conversion factor from power density to power density per unit steradian, see [RL86]).

Our particular task in volume rendering is to determine the radiance $L(\mathbf{x}_e, \omega)$ at the endpoint \mathbf{x}_e of each primary ray of light propagating in direction ω towards the camera through the volume. The total radiance at the endpoint \mathbf{x}_e is the sum of the attenuated “back-light”, and a term due to the light added by the medium at each sample along the ray (attenuated by the transmittance between the sample point \mathbf{x}_n and \mathbf{x}_e):

$$L(\mathbf{x}_e, \omega) = L(\mathbf{x}_s, \omega)T(\mathbf{x}_s, \mathbf{x}_e) + \Delta x \sum_n T(\mathbf{x}_n, \mathbf{x}_e)Q(\mathbf{x}_n). \quad (3)$$

The transmittance function $T(\mathbf{x}_a, \mathbf{x}_b) = \exp(-\tau(\mathbf{x}_a, \mathbf{x}_b))$ gives the attenuation of radiance due to absorption and out-scattering along a given ray segment with endpoints \mathbf{x}_a and \mathbf{x}_b and optical depth (number of mean-free-paths) $\tau = \int_{\mathbf{x}_a}^{\mathbf{x}_b} \sigma_t(\mathbf{x}) d\mathbf{x}$. This is computed numerically by breaking the ray into segments. With segment endpoints \mathbf{x}_n , and segment length Δx , we have in discretized form $T(\mathbf{x}_s, \mathbf{x}_e) \approx \exp(-\sum_n \sigma_t(\mathbf{x}_n)\Delta x)$.

To develop the radiative diffusion approximation, we first separate the total radiance into the sum of the reduced incident light (also referred to as the unscattered light) $L_{ri}(\mathbf{x}, \omega)$ which is the incident radiance from external sources attenuated by extinction, and the medium radiance $L_m(\mathbf{x}, \omega)$, which is generated by both scattering of the reduced incident light and emission from the medium:

$$L(\mathbf{x}, \omega) = L_{ri}(\mathbf{x}, \omega) + L_m(\mathbf{x}, \omega). \quad (4)$$

Substituting Eqn. (4) into Eqn. (2), we get a decomposition of in-scattered light, which results in:

$$4\pi Q(\mathbf{x}) = q_{ri}(\mathbf{x}) + q_m(\mathbf{x}) + j(\mathbf{x}) \quad (5)$$

where

$$\begin{aligned} q_{ri}(\mathbf{x}) &:= \sigma_s \int_{\Omega} L_{ri}(\mathbf{x}, \omega) d\omega, \\ q_m(\mathbf{x}) &:= \sigma_s \int_{\Omega} L_m(\mathbf{x}, \omega) d\omega. \end{aligned} \quad (6)$$

The source field generated by the reduced incident light, q_{ri} , is determined by the choice of external lighting, and we will assume that it has been precomputed. For example in the case of an external incident directional light of radiance L_l in direction ω_l , it has the simple form $q_{ri}(\mathbf{x}) = L_l \sigma_s(\mathbf{x})T(\mathbf{x})$, where $T(\mathbf{x})$ is the transmittance from the point \mathbf{x} in the volume along the ray in the direction towards the light (commonly referred to as a *deep shadow map*). This can be easily generalized to other forms of external lighting.

We can think of the medium radiance $L_m(\mathbf{x}, \omega)$ as the multiply-scattered “diffuse” light (though it includes all orders of scattering, as well as the self-emitted radiance). The diffusion approximation consists of the assumption that we can express the medium radiance as the following expansion in the first two angular moments of the radiation field:

$$L_m(\mathbf{x}, \omega) \approx \frac{1}{4\pi} \phi(\mathbf{x}) + \frac{3}{4\pi} \omega \cdot \mathbf{E}(\mathbf{x}). \quad (7)$$

In this approximation, the diffuse light is describable entirely by two fields: the medium *fluence* (or scalar irradiance, the zeroth angular moment of the medium radiance) $\phi(\mathbf{x}) = \int_{\Omega} L_m(\mathbf{x}, \omega) d\omega$ and the medium *flux* (or vector irradiance, the first angular moment of the medium radiance) $\mathbf{E}(\mathbf{x}) = \int_{\Omega} \omega L_m(\mathbf{x}, \omega) d\omega$. In fact, this is a good approximation in highly opaque media where multiple scattering dominates. This is because multiple scattering tends to reduce the higher-order angular moments [Sta95]. The diffusion approximation provides an efficient means of solving for $\phi(\mathbf{x})$, which determines $q_m(\mathbf{x})$ in Eqn. 6.

The source field expressed in terms of $\phi(\mathbf{x})$ is

$$4\pi Q(\mathbf{x}) = q_{ri}(\mathbf{x}) + \sigma_s(\mathbf{x})\phi(\mathbf{x}) + j(\mathbf{x}) \quad (8)$$

To derive the equation for $\phi(\mathbf{x})$ in the diffusion approximation, we start with the observation that, by definition, the unscattered light satisfies the following restricted form of the RTE excluding the scattering term:

$$(\omega \cdot \nabla)L_{ri}(\mathbf{x}, \omega) = -\sigma_t(\mathbf{x})L_{ri}(\mathbf{x}, \omega). \quad (9)$$

Thus substituting Eqn. (4) into Eqn. (1), and using Eqn. (9), we find that the medium radiance satisfies the RTE:

$$(\omega \cdot \nabla)L_m(\mathbf{x}, \omega) = -\sigma_t(\mathbf{x})L_m(\mathbf{x}, \omega) + Q(\mathbf{x}, \omega). \quad (10)$$

To obtain equations for $\phi(\mathbf{x})$ and $\mathbf{E}(\mathbf{x})$, we substitute Eqn. (7) into Eqn. (10), and integrate over $\int_{\Omega} d\omega$ and $\int_{\Omega} \omega d\omega$ respectively, giving:

$$\begin{aligned} \nabla \cdot \mathbf{E}(\mathbf{x}) &= -\sigma_a(\mathbf{x})\phi(\mathbf{x}) + q_{ri}(\mathbf{x}) + j(\mathbf{x}), \\ \mathbf{E}(\mathbf{x}) &= -\frac{1}{3\sigma_t(\mathbf{x})} \nabla \phi(\mathbf{x}) = -D(\mathbf{x}) \nabla \phi(\mathbf{x}). \end{aligned} \quad (11)$$

The first line of Eqn. (11) is exact and (when integrated over volume) simply states conservation of energy, i.e., that in a steady-state solution of the RTE the total power of diffuse light emitted from any given region must equal the rate of increase in diffuse




Figure 2: Schematic of the volume rendering process. The unscattered attenuated incident light L_{ri} generates the source field $q_{ri}(\mathbf{x})$, and emission within the volume generates the source field $j(\mathbf{x})$. The multiply-scattered light $\phi(\mathbf{x})$ is then determined by solving the diffusion Eqn. (12) with the source fields $q_{ri}(\mathbf{x})$ and $j(\mathbf{x})$. The fields $q_{ri}(\mathbf{x})$, $j(\mathbf{x})$, and $\phi(\mathbf{x})$ are then used to compute the source term $Q(\mathbf{x}_n)$ at each ray segment according to Eqn. (8), which is accumulated along the ray to compute the observed radiance according to Eqn. (3).

radiant energy in the region due to in-scattering and emission, minus the rate of decrease due to absorption. The second line of Eqn. (11), commonly termed Fick's first law of diffusion, is the statement that in the isotropic diffusion approximation the flux \mathbf{E} is proportional to the gradient of ϕ , where $D(\mathbf{x}) := [3\sigma_t(\mathbf{x})]^{-1}$ is the *classical diffusion coefficient*, which has dimensions of length. Flux-limited diffusion theory, to be introduced in the next section, amounts to a modification of this diffusion coefficient.

Substituting Fick's first law into the first line of Eqn. (11) yields Fick's second law, the radiative diffusion equation:

$$\nabla \cdot (D(\mathbf{x}) \nabla \phi(\mathbf{x})) = \sigma_a(\mathbf{x}) \phi(\mathbf{x}) - q_{ri}(\mathbf{x}) - j(\mathbf{x}). \quad (12)$$

This describes how the zeroth moment of the multiply-scattered radiance (or fluence), $\phi(\mathbf{x})$, is determined by the incident unscattered light and the emission, represented by the source terms $q_{ri}(\mathbf{x})$ and $j(\mathbf{x})$. This is a partial differential equation defining a boundary value problem which can be solved for $\phi(\mathbf{x})$ in the volume domain, given specified conditions on $\phi(\mathbf{x})$ to impose at the domain boundary. We use Dirichlet boundary conditions where $\phi(\mathbf{x})$ is simply set to zero on the boundary. Though this is physically inaccurate, our procedure of using the unscattered light to seed the diffusion solve in the interior of the volume minimizes the influence of the boundaries.

Volume rendering an image is then completed by marching through the volume along the primary ray in order to compute the sum in Eqn. (3), using the computed fluence field $\phi(\mathbf{x})$. The whole process is summarized in Fig. 2.

4. Flux-Limited Diffusion

In this section, we introduce the theoretical ideas of flux-limited diffusion theory (FLD), an enhancement to the classical diffusion theory which proves to be more accurate, especially in heterogeneous media with low opacity regions. In outline, the method amounts to replacing the diffusion coefficient $D(\mathbf{x}) = [3\sigma_t(\mathbf{x})]^{-1}$ with an expression which depends on the fluence $\phi(\mathbf{x})$ and the magnitude of its gradient, $|\nabla \phi(\mathbf{x})|$. FLD was developed originally

for radiative transfer applications in astrophysics [LP81, TS01] and nuclear physics [LK74]. Here, we will state the mathematical framework of FLD required for its application, but without giving the full mathematical derivation, aiming instead to give intuition for its physical meaning.

The classical diffusion approximation cannot be applied straightforwardly to the computation of multiple scattering in heterogeneous media, as it breaks down and violates the true RTE in regions of low extinction σ_t (or equivalently large mean free path). The first indication of this is that in the limiting case of a vacuum (i.e., infinite mean-free-path), the radiative diffusion equation, Eqn. (12), becomes singular, because the diffusion coefficient $D(\mathbf{x})$ diverges in the limit $\sigma_t \rightarrow 0$. A more rigorous demonstration that this is inconsistent with the true RTE is as follows. The medium radiance is non-negative and therefore manifestly satisfies the following inequality, where $\hat{\mathbf{n}}$ is any unit vector:

$$\int_{\Omega} (1 - \hat{\mathbf{n}} \cdot \boldsymbol{\omega}) L_m(\mathbf{x}, \boldsymbol{\omega}) d\boldsymbol{\omega} \geq 0. \quad (13)$$

This reduces to the statement that $\phi(\mathbf{x}) \geq \hat{\mathbf{n}} \cdot \mathbf{E}(\mathbf{x})$ for all unit vectors $\hat{\mathbf{n}}$, which implies the constraint

$$|\mathbf{E}(\mathbf{x})| \leq \phi(\mathbf{x}). \quad (14)$$

This states that at any point in the radiation field the magnitude of the flux does not exceed the fluence. Physically this is related to the fact that radiation propagates at the finite speed of light, so energy cannot be transported at an arbitrarily high rate from point to point (thus this flux constraint is sometimes said to enforce "causality").

In CDA, $\phi(\mathbf{x})$ satisfies Eqn. (12) and the corresponding flux is given by the Fick's law relationship (the second line of Eqn. (11)). However, there is no mechanism to prevent a solution where $|\mathbf{E}(\mathbf{x})| = |D(\mathbf{x}) \nabla \phi(\mathbf{x})|$ is arbitrarily large compared to $\phi(\mathbf{x})$, which violates the constraint in Eqn. (14). This suggests that CDA can be improved by explicitly enforcing the constraint, which is the core idea of FLD theory.

To achieve this, in FLD theory the diffusion coefficient in Fick's law is modified to

$$D_F(\mathbf{x}) = \frac{\mathcal{F}(\mathcal{R}(\mathbf{x}))}{\sigma_t(\mathbf{x})}, \quad (15)$$

where the dimensionless $\mathcal{F}(\cdot)$ is termed the *flux limiter*. CDA corresponds to setting \mathcal{F} to the constant 1/3. In FLD, \mathcal{F} is a function of the dimensionless variable $\mathcal{R}(\mathbf{x})$, termed the *Knudsen number*, which is defined as follows:

$$\mathcal{R}(\mathbf{x}) := \frac{|\nabla \phi(\mathbf{x})|}{\sigma_t(\mathbf{x}) \phi(\mathbf{x})}. \quad (16)$$

Via the modified Fick's first law $\mathbf{E}(\mathbf{x}) = -D_F(\mathbf{x}) \nabla \phi(\mathbf{x})$, the flux-limiting constraint Eqn. (14) expressed in terms of the flux limiter and Knudsen number becomes:

$$\mathcal{F} \leq \frac{1}{\mathcal{R}}. \quad (17)$$

The physical meaning of the Knudsen number \mathcal{R} can be understood roughly as:

$$\mathcal{R} \sim \frac{\text{mean free path}}{\text{length scale of } \phi \text{ variations}}. \quad (18)$$

When the local mean free path is short relative to the local length scale of changes in the fluence, $\mathcal{R} \ll 1$ and the radiation field is in a "diffusive" regime where CDA, $\mathcal{F} = 1/3$, is a good approximation.

Table 2: Flux limiters, $\mathcal{F}(\mathcal{R})$.

Flux limiter	$\mathcal{F}(\mathcal{R})$
sum [BW82]	$(3 + \mathcal{R})^{-1}$
max [BW82]	$\max(3, \mathcal{R})^{-1}$
Kershaw [Ker76]	$2(3 + \sqrt{9 + 4\mathcal{R}^2})^{-1}$
Larsen- n [LK74]	$(3^n + \mathcal{R}^n)^{-\frac{1}{n}}$
Levermore–Pomraning [LP81]	$\frac{1}{\mathcal{R}} \left(\coth(\mathcal{R}) - \frac{1}{\mathcal{R}} \right)$

In this case Eqn. (17) is satisfied automatically. However, when the local mean free path is much longer than this length scale, $\mathcal{R} \gg 1$ and the radiation field enters the “transport” regime where interaction events are rare over the length scale being considered. Thus in FLD, in order to satisfy the constraint Eqn. (17), we use a modified flux limiter which satisfies the limits:

$$\lim_{\mathcal{R} \rightarrow 0} \mathcal{F}(\mathcal{R}) = \frac{1}{3}, \quad \lim_{\mathcal{R} \rightarrow \infty} \mathcal{R} \mathcal{F}(\mathcal{R}) = 1. \quad (19)$$

A variety of methods have been developed to find specific functional forms for $\mathcal{F}(\mathcal{R})$ satisfying the constraints in Eqn. (19). Some approaches use forms which simply satisfy the desired limits [LK74, BW82]. Others derived flux limiters based on certain assumptions about the smoothness of the moments [LP81, Ker76]. We give the expressions for these simple widely used flux limiters in Table 2. Other strongly related work includes Minerbo’s [Min78] flux limiters which are based on statistical mechanics of the radiation field, and a method termed “M1-theory” which works by making assumptions about the local rotational symmetry of the radiation field and deriving the corresponding pressure tensor closure [OAH00]. We do not expand on either of these as this would exceed the scope of this paper.

Regardless of the specific choice of flux limiter, the resulting flux-limited radiative diffusion equation has the form

$$\nabla \cdot (D_F \nabla \phi(\mathbf{x})) = \sigma_a(\mathbf{x})\phi(\mathbf{x}) - q_{ri}(\mathbf{x}) - j(\mathbf{x}). \quad (20)$$

Here D_F implicitly depends on $\phi(\mathbf{x})$ via Eqn. (16), which makes this a non-linear PDE, unlike the classical diffusion Eqn. (12) which is linear. As in the CDA case, we use Dirichlet boundary conditions to specify $\phi(\mathbf{x})$ on the domain edge. In summary, FLD takes the theoretical framework of CDA and modifies it straightforwardly by replacing the diffusion coefficient with a $\phi(\mathbf{x})$ -dependent expression. In Sec. 5, we give the details of implementing a numerical CDA solver and extending it to support the FLD technique.

Note that in the limit $\mathcal{R} \rightarrow \infty$, the divergence term in the diffusion equation becomes $-\nabla \cdot (\phi(\mathbf{x})\hat{\mathbf{n}})$, which effectively transforms the diffusion equation into an advection equation for the fluence. This also explains why the flux limiter is chosen to tend to exactly $1/\mathcal{R}$ in Eqn. (19), which is a stronger requirement than imposed by the constraint of Eqn. (17): FLD not only suppresses the flux in the transport regime, it also saturates it at the appropriate value to ensure correct free propagation (at the level of the approximation).

A deeper examination of the derivation of CDA gives further insight into the physical content of FLD. As discussed in Section 3, CDA can be derived by expressing the radiance in terms of its first two moments only, leading to the diffusion equation which determines these two moments. It was not specified what assumptions we made about the higher moments—in general these mo-

ments are non-zero as they are coupled to the first and second moments via a hierarchy. In fact, a deeper description of the diffusion approximation is that the moment hierarchy is closed by choosing a specific form for the second moment, termed the *radiation pressure tensor*, $P_{ij}(\mathbf{x}) = \int_{\Omega} \omega_i \omega_j L_m(\mathbf{x}, \omega) d\omega$. In the case of CDA, the radiation pressure tensor is approximated as isotropic, $P_{ij} = \delta_{ij}\phi/3$, which leads to the approximate two-term expansion of the radiance given in Eqn. (7).

In the case of FLD, the form of the pressure tensor is different. In terms of the *Eddington tensor* E_{ij} , defined by $P_{ij} = E_{ij}\phi$, in CDA $E_{ij} = \delta_{ij}/3$, while in FLD

$$E_{ij} = \frac{1}{2}(1 - \mathcal{X})\delta_{ij} + \frac{1}{2}(3\mathcal{X} - 1)\hat{\mathbf{n}}_i\hat{\mathbf{n}}_j, \quad (21)$$

where $\hat{\mathbf{n}}$ is the unit vector in the direction of the negative fluence gradient, $\hat{\mathbf{n}}(\mathbf{x}) = -\nabla\phi(\mathbf{x})/|\nabla\phi(\mathbf{x})|$, and \mathcal{X} is a dimensionless function called the *Eddington factor*. In a full analysis it has been shown [Pom82, Lev84] that the following approximate relationship holds between the flux limiter \mathcal{F} and the Eddington factor \mathcal{X} :

$$\mathcal{X} = \mathcal{F} + \mathcal{F}^2\mathcal{R}^2. \quad (22)$$

This implies that in the diffusive and transport limits respectively, the FLD Eddington tensor becomes

$$\lim_{\mathcal{R} \rightarrow 0} E_{ij} = \delta_{ij}/3, \quad \lim_{\mathcal{R} \rightarrow \infty} E_{ij} = \hat{\mathbf{n}}_i\hat{\mathbf{n}}_j. \quad (23)$$

From this we see explicitly that in the transport regime, FLD produces a radiation pressure tensor closure which corresponds to free propagation of rays down the local energy density gradient. It should be borne in mind however that though FLD deals in a more physically sensible manner with the transition to the transport regime than CDA, it is still just an approximation to the full RTE, as all but the first few moments of the radiation field are used.

4.1. Point Source in Homogeneous Medium

In this section we apply the ideas introduced above to the problem of a unit power point light embedded in an infinite homogeneous scattering medium with extinction coefficient σ_t and albedo α . This is a spherically symmetric problem for which the exact solution is known to both the full RTE and CDA. We use this example both to illustrate the general idea of FLD and as a testbed for our numerical solver.

The exact solution of the RTE for the fluence as a function of radial distance from the point source, and an excellent approximation to it, termed the Grosjean solution [Gro58], are quoted in previous work [D111]. Parameterizing the radial distance r from the point source in terms of the dimensionless *optical depth* $\tau = \sigma_t r$, the Grosjean solution for the fluence is given (in the case of an isotropic phase function) by:

$$\phi(\tau) = \frac{\sigma_t^2}{4\pi} \left(\frac{e^{-\tau}}{\tau^2} + \frac{3\alpha}{2-\alpha} \frac{e^{-\lambda\tau}}{\tau} \right), \quad (24)$$

where $\lambda^2 := \frac{3(1-\alpha)}{2-\alpha}$. Note that the σ_t^2 factor occurs since ϕ has units of radiance (power per unit area), and we used a unit power source. The expression inside the brackets is dimensionless. In Fig. 3(a), we plot the “normalized fluence”

$$\tilde{\phi}(\tau) := 4\pi\phi(\tau)/\sigma_t^2. \quad (25)$$




Figure 3: Analytical “universal” solutions for normalized fluence versus radius for a unit power point source embedded in a homogeneous medium. We plot in (a) the Grosjean transport solution Eqn. (24) and in (b) the classical diffusion approximation Eqn. (27), both for a range of albedos between 0 and 1. To show the solutions for all possible σ_t we plot normalized fluence (defined in Eqn. (25)) as a function of optical depth $\tau = \sigma_t r$. In (c) we plot the Knudsen number \mathcal{R} (Eqn. (16)) of the Grosjean solution.

as a function of τ . These figures show the σ_t -independent “universal” solution to the point source scattering problem.

The CDA in Eqn. (12) reduces in this homogeneous emission-only case to a screened Poisson equation:

$$D\nabla^2\phi(\mathbf{x}) = \sigma_a(\mathbf{x})\phi(\mathbf{x}) - j(\mathbf{x}). \quad (26)$$

The solution to this, for a unit power point source emission $j(\mathbf{x}) = \delta(\mathbf{x})$, known as the CDA Greens function [JMLH01], is given by

$$\phi(\tau) = \frac{3\sigma_t^2}{4\pi} \frac{e^{-\sqrt{3(1-\alpha)}\tau}}{\tau}. \quad (27)$$

Fig. 3(b) plots the CDA Greens function as a function of optical depth. As it has been discussed in detail [DI11], we see that CDA fails to generate the τ^{-2} power law term in the Grosjean solution, and also has stronger exponential suppression of the τ^{-1} diffusion power law at large radii, which causes it to underestimate the fluence.

In the Grosjean solution the Knudsen number $\mathcal{R} \rightarrow 2/\tau$ as $\tau \rightarrow 0$ (see Fig. 3(c)). Thus, sufficiently close to the point source, i.e., roughly within a mean-free-path, the radiation field is in the transport regime according to the criterion described in Section 4. In this region, the inverse square term in Eqn. (24) dominates (and in the $\alpha = 0$ limit it is the exact transport solution, as can be verified by trivial application of the RTE without scattering). In the terminology of [DI11], this term represents the *ballistic fluence* due to light propagating unscattered from the source.

As pointed out by Levermore and Pomraning [LP81], in spherical symmetry the FLD Eqn. (20) with a point emission source becomes in the transport regime (according to the $\mathcal{R} \rightarrow \infty$ limiting form in Eqn. (19)) the following linear first-order ordinary differential equation:

$$\frac{1}{\tau^2} \frac{d}{d\tau} (\tau^2 \phi(\tau)) + (1 - \alpha)\phi(\tau) = j(r)/\sigma_t, \quad (28)$$

whose corresponding solution for a delta function j is

$$\phi(\tau) = \frac{\sigma_t^2}{4\pi} \frac{e^{-(1-\alpha)\tau}}{\tau^2}, \quad (29)$$

which tends to the ballistic term in the Grosjean solution as $\tau \rightarrow 0$. This verifies that FLD correctly reproduces the transport limit for

the point source problem. Thus, one reasonable test of our formulation is to check that FLD gives this inverse square behavior of the fluence near an emitting point source. However, note that we do not expect FLD to perfectly match the transport theory throughout the domain, as the FLD equation is merely an approximation to the RTE.

5. Numerical Method

We provide here a numerical method for solving Eqn. (20) to obtain the multiply-scattered fluence $\phi(\mathbf{x})$. We assume a uniform 3D grid with voxel centers indexed by integer coordinates $p = (i, j, k)$, and voxel edge length Δl . For notational convenience, we absorb the external illumination term $q_{ri}(\mathbf{x})$ into the emission term $j(\mathbf{x})$. Hence, the voxelized input fields are the emission field j_p (units of power per unit volume), the extinction field $\sigma_p := \sigma_t|_{ijk}$ (units of inverse length), and the albedo field α_p (dimensionless). For brevity we write the flux-limited diffusion coefficient $D_F|_{ijk}$ as D_p (units of length).

Discretizing the LHS of Eqn. (20) at each grid point, the x-derivative in the divergence term reads

$$\frac{D_{i+1/2,j,k}(\phi_{i+1,j,k} - \phi_p) - D_{i-1/2,j,k}(\phi_p - \phi_{i-1,j,k})}{\Delta l^2}, \quad (30)$$

where $D_{i+1/2,j,k}$ denotes the diffusion coefficient interpolated linearly at $(i + 1/2, j, k)$. Thus, defining $s(p)$ as the 6-point stencil surrounding point p , so $s(p)$ ranges over the six points $(i \pm 1, j \pm 1, k \pm 1)$, we obtain the following discretized form of Eqn. (20):

$$\frac{\sum_s D_{ps} \phi_s - \phi_p \sum_s D_{ps}}{\Delta l^2} = (1 - \alpha_p) \sigma_p \phi_p - j_p, \quad (31)$$

where we abbreviated $D_{ps} := (D_p + D_s)/2$. Given the current set of ϕ_p values on the grid, we can estimate the accuracy of the solution by the *residual* R_p , which is defined as the LHS minus the RHS of Eqn. (31). Our goal is to solve for a set of values of ϕ_p at every interior grid point which minimizes the root mean square (RMS) residual. However, in order to apply the FLD method, D_p at each point must also satisfy Eqn. (15), which makes Eqn. (31) a non-linear equation involving the values of ϕ and its partial derivatives on the stencil (the specific form of this equation depends on the choice of the flux limiter \mathcal{F}).




Figure 4: Convergence behavior of our FLD solver. In (a) we show the normalized RMS residual \bar{R}/\bar{j} as a function of iterations for the problem described in Section 5, run in double precision, with a 51^3 grid. In (b) we show the number of iterations required to reach the $\bar{R}/\bar{j} = 10^{-6}$ convergence criterion as a function of grid resolution for the same test case. In (c) we demonstrate the effect of extinction tolerance σ_ϵ on convergence of the Nebulae dataset in Fig. 1, here in single precision.

To find a consistent solution for both ϕ_p and D_p , we employ a Gauss-Seidel relaxation scheme. We solve Eqn. (31) for ϕ_p to obtain the following local update rule for ϕ_p :

$$\phi_p \leftarrow \frac{j_p \Delta l^2 + \sum_s D_{ps} \phi_s}{(1 - \alpha_p) \sigma_p \Delta l^2 + \sum_s D_{ps}}. \quad (32)$$

We also update D_p at each stencil according to

$$D_p \leftarrow \frac{\mathcal{F}(\mathcal{R}_p)}{\max(\sigma_p, \sigma_\epsilon)}. \quad (33)$$

The threshold σ_ϵ introduced here is discussed later in this section. The local Knudsen number \mathcal{R}_p is computed using the gradient vector at p , which is discretized as

$$\nabla \phi_p = \frac{1}{2\Delta l} \begin{pmatrix} \phi_{i+1,j,k} - \phi_{i-1,j,k} \\ \phi_{i,j+1,k} - \phi_{i,j-1,k} \\ \phi_{i,j,k+1} - \phi_{i,j,k-1} \end{pmatrix}. \quad (34)$$

Algorithm 1: Gauss-Seidel FLD solver

Input voxel grids j_p , σ_p , α_p of extent L , voxel size Δl
 Impose extinction tolerance $\sigma_p \leftarrow \max(\sigma_p, \sigma_\epsilon)$
 Compute RMS emission field \bar{j}
 Initialize voxel grids $\phi_p = \epsilon \bar{j} \Delta l$, $D_p = \epsilon \Delta l$, $\forall p$
 Choose flux limiter $\mathcal{F}(\mathcal{R})$, from Table 2
 Choose SOR factor, $0 < \omega < 2$
repeat
 parallelfor $\forall p$ (non-boundary red points) **do**
 Compute \mathcal{R} via Eqn. (35)
 Update D_p via Eqn. (33), using $\mathcal{F}(\mathcal{R})$
 Compute $D_{ps} = (D_p + D_s)/2 \forall 6$ points s
 Compute ϕ'_p via Eqn. (32), using D_{ps}
 Update $\phi_p \leftarrow \omega \phi'_p + (1 - \omega) \phi_p$, Eqn. (36)
 Compute \mathcal{R}_p via Eqn. (37)
 end
 parallelfor $\forall p$ (non-boundary black points) **do**
 As above.
 end
 Compute RMS of residual field, \bar{R}
until $\bar{R} < 10^{-6} \bar{j}$

We then compute

$$\mathcal{R}_p = \frac{\max(|\nabla \phi_p|, \epsilon \bar{j})}{\max(\sigma_p \phi_p, \epsilon \bar{j})}. \quad (35)$$

This sets a lower bound $\epsilon \bar{j}$ on $|\nabla \phi_p|$ and $\sigma_p \phi_p$, where $\epsilon = 10^{-20}$ and \bar{j} is the RMS value of the emission field, to deal with the case when either ϕ_p or $|\nabla \phi_p|$ is extremely small.

For parallelization we use a Red-Black Gauss-Seidel update process. We color the grid points as a red-black checkerboard, then apply the updates of Eqn. (32) and (33) concurrently at all red points (one “pass”) followed by a pass over all black points. Each stencil computes and updates only its own local ϕ_p and D_p values during the red (black) pass, which is then stored in the grid for use in the following black (red) pass. The two passes constitute one iteration. The ϕ_p and D_p grids are initialized to small tolerances $\phi_p = \epsilon \bar{j} \Delta l$ and $D_p = \epsilon \Delta l$. Initially, these grid values do not satisfy Eqn. (32) and Eqn. (33), but over a number of iterations they converge to a consistent solution of both equations over the whole grid. At the grid boundaries, we impose Dirichlet boundary conditions with $\phi = \epsilon \bar{j} \Delta l$ by ignoring the voxels abutting the six grid faces during the update. This introduces some unphysical effects near the boundaries, but does not significantly corrupt the solution in the bulk of the grid.

To improve the rate of convergence, we use the successive over-relaxation (SOR) technique. Defining the over-relaxation parameter ω , where $0 < \omega < 2$, the update of ϕ_p at each stencil is modified to

$$\phi_p \leftarrow \omega \phi'_p + (1 - \omega) \phi_p \quad (36)$$

where ϕ'_p is the updated value from Eqn. (32). To monitor convergence, we compute a normalized residual as the RMS of the residual field \bar{R} (ignoring voxels close to the boundary, to avoid boundary effects), divided by \bar{j} . Writing Eqn. (32) as $\phi_p \leftarrow \text{numerator}/\text{denominator}$, we compute \mathcal{R}_p as:

$$\mathcal{R}_p = (\text{numerator} - \phi_p \cdot \text{denominator}) / \Delta l^2. \quad (37)$$

When \bar{R}/\bar{j} falls below a suitable tolerance, e.g. 10^{-6} for double precision calculations, we regard the numerical solution as having converged. An overview of the complete process is given in Algorithm 1.




Figure 5: Numerical solution for point source normalized fluence $\tilde{\phi}$ (defined in Eqn. 25) for various albedos. In (a) and (b) we used the CDA solver, and in (c) we used our FLD solver. In all figures, points indicate the numerical solution for $\tilde{\phi}$ at given optical depth τ , and solid/dotted curves show the CDA theory Eqn. (27) and Grosjean transport theory Eqn. (24) as labelled. In (a) we show the numerical solution of CDA with Dirichlet boundary conditions set from the analytical solution. It closely matches the CDA theory, demonstrating that the deviation of the numerical solve from the theory in (b) is entirely due to the effect of the $\phi = 0$ Dirichlet boundary condition applied in this case. In (c) we use the $\phi = 0$ boundary condition.

In Eqn. (33), we impose a lower bound on the extinction σ_e , as we found empirically that this tolerance has an effect on the convergence rate in the presence of vacuum regions. This effect is shown in Fig. 4(c), using our GPU solver applied to the Nebulae dataset. A value of $\sigma_e = 10^{-3}/L$ (where L is the grid extent), corresponding to a maximum optical depth of 1000 box lengths, was sufficient to ensure good convergence with all of our test datasets.

We implemented the numerical method described in both a tested solver in double precision on the CPU, and a high performance GPU solver implemented in single precision via the NVIDIA CUDA framework. In Figures 4(a) and 4(b) we show the convergence properties of our CPU FLD solver on a simple test case where the extinction field defines an embedded spherical cloud (with a vacuum exterior) and an internal emission field generated via Perlin noise. These figures demonstrate that our numerical method converges over a broad range of SOR parameters.

6. Results

In this section we apply our numerical FLD solver first to the point source problem discussed in Section 4.1, and then to various heterogeneous datasets.

6.1. Point Source in Homogeneous Medium

We first applied our FLD solver to the point source problem presented in Section 4.1 to verify that it behaves according to the theory. The point source was modeled as a single emitting voxel at the center of the numerical grid, which had a resolution of 127^3 , and a homogeneous extinction corresponding to $\tau = 4$ across the width of the grid. We set j within the emitting voxel to the reciprocal of the voxel volume in order to correctly discretize a unit power point light. We ran both the CDA and FLD solvers (where the CDA solver is simply the FLD solver with flux limiter set to constant $\mathcal{F} = 1/3$) for a range of albedos, and then extracted the normalized fluence as a function of radial distance.

In Fig. 5, we compare our numerical results with the analytical solutions. Fig. 5(b) shows the solution obtained using CDA, and Fig. 5(c) shows the solutions obtained via FLD using the

Levermore–Pomraning flux limiter. With the CDA solver, at low to medium albedo we obtain a good match between classical diffusion theory and the numerics, as expected since the CDA numerical solver is based on CDA theory. At high albedo, the effect of the boundaries becomes more pronounced, making the obtained CDA solution deviate from the CDA theory of the case of an infinite homogeneous medium. However, we do not expect a perfect match between the FLD numerical solver and the Grosjean theory, since FLD is merely an approximation to the full transport theory. Despite this, with the FLD solver we obtain a rough match to the Grosjean “ballistic” behavior at low optical depth, and the Levermore–Pomraning flux limiter seems to produce a reasonable quality match across the whole albedo range (the other flux limiters produce a similar though slightly worse match). This indicates that the FLD approximation of the transition to the transport regime is working reasonably well.

6.2. Practical Examples

In this section, we compare renderings of heterogeneous datasets obtained via FLD, CDA, and a volumetric path tracer which serves as ground truth. The path tracer uses the standard Woodcock tracking method [RSK06] and was validated using Chandrasekhar’s analytical solution for multiple scattering in a slab geometry [Cha60].

In all cases we illuminated the volumes with a single directional external light source L_l . For the FLD and CDA tests, the external lighting produces the multiple scattering source term $q_{ri}(\mathbf{x}) = L_l \sigma_s(\mathbf{x}) T(\mathbf{x})$, which we computed via raymarching and stored at the full resolution of the dataset extinction field. We solve Eqn. (20) for the fluence $\phi(\mathbf{x})$ using Algorithm 1, with source term $q_{ri}(\mathbf{x})$. We use the Levermore–Pomraning flux-limiter in all examples (the other flux-limiters were tried and produced very similar results). Then via Eqn. (8) we complete the primary raymarch Eqn. (3) to render the image. Since multiply-scattered light tends to reduce higher-order angular moments, $\phi(\mathbf{x})$ has a much lower spatial frequency than the extinction field and therefore resolving it at the same resolution is not necessary. In fact, a major benefit of the diffusion approach is the ability to decouple the computation of the multiply-scattered light from the singly-scattered light,




Figure 7: Nebulae dataset rendered for albedo values from 0.1 (left) to 1.0 (right). The path-traced reference (top row) is compared against FLD (middle row) and CDA (bottom row).

and solve for the multiply-scattered contribution more efficiently. To exploit this we therefore, in all our examples, solved for $\phi(\mathbf{x})$ on grids with a quarter of their original resolution. Fig. 6 demonstrates that the effect of solving for the multiply-scattered light on such a coarser grid has a low visual impact.

We implemented all three approaches in single precision using the NVIDIA CUDA framework. When measuring performance, we included in the CDA and FLD cases the time taken for primary raymarching and generation of the source term (light baking), as well as the time taken during the diffusion solve. All results were obtained with a NVIDIA GTX TITAN GPU. We ran the diffusion solver until convergence at the level of $\bar{R}/\bar{j} = 10^{-4}$, which is the practical limit for the single precision computation. With all our path-traced images, we used 1000 Monte Carlo samples per-pixel, which we found to strike a good balance between image quality and computation time.


All images were rendered at a resolution of 1024^2 pixels, and tone-mapped using standard gamma compression. For each dataset, exposure was adjusted for ground truth rendering to avoid blown highlights and then the same tone mapping was used for FLD and CDA.

The Nebulae example, shown in Fig. 1, was generated from a noise-displaced implicit surface. Color is introduced by a color-dependent scaling of the extinction field. The resolution of the

dataset is 200^3 and the albedo is 0.9. Note the differences of the shadowed regions and the volume blob at the bottom of the dataset, which only receives indirect light coming from other illuminated parts of the volume. While CDA is completely black in those regions, FLD finds a solution which is much closer to the path-traced result. To compare the behavior of both methods as albedo is varied, we show in Fig. 7 that the rendering of the Nebulae dataset with FLD produces better results than CDA over the whole range of albedos, though with a more pronounced difference at high albedos.

The example in Fig. 8 is based on a CT-scan dataset, which has a resolution of $256 \times 329 \times 256$. In this example, the density field in the original dataset is mapped to a per-channel albedo field and a (color channel independent) extinction field via a transfer function which colors tissue and bone areas distinctly. The image shows that FLD better captures indirect light at the neck and the translucent nature of tissue in the path-traced rendering, when compared to CDA.

The example in Fig. 9 is based on a fluid simulation with a resolution of $201 \times 272 \times 230$. There is an external light source and a constant albedo of 1.0. In addition, there is a black body emission field which is calculated from the temperature field of the simulation. Again FLD matches the ground truth closer than CDA.



(a) Original resolution (b) Quarter resolution

Figure 6: Comparison of our method when solving for multiply-scattered light using FLD at original resolution 200^3 (left) and at resolution 50^3 (right).

These examples demonstrate that FLD produces results which are closer to path tracing than CDA for a variety of datasets and parameterizations. Significant improvements over CDA appear especially in situations with indirect lighting or transitions from highly dense media to regions of low density or vacuum. As a modified diffusion theory, FLD of course still deviates to a certain extent from the true RTE solution, as all diffusion methods throw away all but the lowest angular moments of the radiation field. Our performance measurements indicate that FLD is roughly twice as computationally expensive as CDA, but is still very fast when compared with path tracing. We noted that in our GPU implementation, this loss in performance is dominated by the additional memory access per voxel which is required to compute the Knudsen number.




Figure 8: MANIX CT-scan dataset with spatially varying albedo. Extinction coefficient and albedo are driven by transfer functions. Note how FLD better captures indirect light at the neck in comparison to CDA.




Figure 9: Fluid simulation dataset lit by internal blackbody emission and an external light source.

7. Conclusion

We introduced the flux-limited diffusion approximation for multiple scattering in participating media. Flux-limited diffusion originated in the field of astrophysics as an improvement to CDA which produces results of higher accuracy when applied to heterogeneous media containing regions of low extinction or vacuum. While FLD is slightly more computationally expensive than CDA, it is still much faster than path tracing. We provided a numerical method which we checked for convergence and validated using an analytical solution for scattering of a point source. We rendered various heterogeneous volumetric datasets with FLD, CDA, and path tracing, and demonstrated that FLD produces a better match to path-traced ground truth than CDA.

For future work, we suggest that the extension of FLD to anisotropic phase functions may be of particular importance. In the anisotropic radiative diffusion theory, the extinction is replaced by a “reduced extinction”, $\sigma'_t = \sigma_t(1 - \alpha g)$, where $g \in [0, 1]$ is the phase function anisotropy factor [Sta95, JMLH01]. It causes a divergence in the diffusion coefficient, and correspondingly the flux, as $g \rightarrow 1$ (strong forward scattering) and $\alpha \rightarrow 1$ (high albedo), even in optically thick media. Thus, in the case of strongly forward scattering high-albedo media, flux-limiting is an essential ingredient to obtain physically accurate results from diffusion meth-

ods. This is of practical importance in the application of diffusion-based methods to, for example, the rendering of Mie scattering [BNM*08, YIC*10] in atmospheric clouds. The form of the flux is also different from Eqn. (11) in the presence of anisotropy, with an additional term involving the first moment of the external flux [Sta95, MSM*04], which suggests that a detailed theoretical analysis is needed to understand how to achieve flux limiting in this case.

There might also be substantial performance benefit to be gained by developing a more sophisticated numerical algorithm for solving the flux-limited diffusion equation than the one presented here, such as a non-linear multigrid solver [BHM00]. Finally, it is worth noting that any existing graphics technique which is based on classical diffusion theory, such as the work of Arbree et al. [AWB11], could be easily augmented with the FLD technique for improved accuracy.

8. Acknowledgements

This work was partially funded by the graduate scholarship program Digital Media, granted by the state Baden-Württemberg, Germany, and the Cluster of Excellence in Simulation Technology (EXC 310/1), granted by Deutsche Forschungsgemeinschaft

(DFG). The authors would like to thank Jose Esteve and Sebastian Herholz for helpful suggestions.

References

- [AWB11] ARBREE A., WALTER B., BALA K.: Heterogeneous subsurface scattering using the finite element method. *IEEE Transactions on Visualization and Computer Graphics* 17, 7 (2011), 956–969. 1, 2, 10
- [BCK*11] BØRLUM J., CHRISTENSEN B. B., KJELDSEN T. K., MIKKELSEN P. T., NOE K. O., RIMESTAD J., MOSEGAARD J.: SSLPV:subsurface light propagation volumes. In *Proc. Symposium on High Performance Graphics* (2011), pp. 7–14. 2
- [BHM00] BRIGGS W. L., HENSON V. E., MCCORMICK S. F.: *A multi-grid tutorial (2nd ed.)*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000. 10
- [BNM*08] BOUTHORS A., NEYRET F., MAX N., BRUNETON E., CRASSIN C.: Interactive multiple anisotropic scattering in clouds. In *Proc. Symposium on Interactive 3D Graphics and Games* (2008), pp. 173–182. 2, 10
- [BW82] BOWERS R. L., WILSON J. R.: A numerical model for stellar core collapse calculations. *Astrophysical Journal Supplement Series* 50 (1982), 115–159. 5
- [Cha60] CHANDRASEKHAR S.: *Radiative transfer*. Dover Books on Intermediate and Advanced Mathematics. 1960. 2, 8
- [CPCP*05] CEREZO E., PÉREZ-CAZORLA F., PUEYO X., SERON F. J., SILLION F.: A survey on participating media rendering techniques. *The Visual Computer* 21, 5 (2005), 303–328. 2
- [DI11] D’EON E., IRVING G.: A quantized-diffusion model for rendering translucent materials. *ACM Transactions on Graphics* 30, 56 (2011), 1–14. 2, 5, 6
- [Fat09] FATTAL R.: Participating media illumination using light propagation maps. *ACM Transactions on Graphics* 28, 1 (2009), 7–18. 2
- [Gro58] GROSJEAN C.: Multiple isotropic scattering in convex homogeneous media bounded by vacuum. In *Proc. of the Second International Conference on the Peaceful Uses of Atomic Energy* (1958), vol. 16, p. 413. 5
- [GRWS04] GEIST R., RASCHE K., WESTALL J., SCHALKOFF R.: Lattice-boltzmann lighting. In *Proc. Symposium on Rendering* (2004), pp. 355–362. 2
- [HCJ13] HABEL R., CHRISTENSEN P. H., JAROSZ W.: Photon beam diffusion: A hybrid monte carlo method for subsurface scattering. *Computer Graphics Forum* 32, 4 (2013). 2
- [Ish78] ISHIMARU A.: *Wave propagation and scattering in random media*, vol. 2. Academic Press, 1978. 2
- [JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proc. 28th Annual Conference on Computer Graphics and Interactive Techniques* (2001), pp. 511–518. 1, 2, 6, 10
- [KD10] KAPLANYAN A., DACHSBACHER C.: Cascaded light propagation volumes for real-time indirect illumination. In *Proc. Symposium on Interactive 3D Graphics and Games* (2010), pp. 99–107. 2
- [Ker76] KERSHAW D. S.: Flux limiting nature’s own way - a new method for numerical solution of the transport equation. *Lawrence Livermore National Laboratory Report UCRL-78378* (1976). 5
- [KVH84] KAJIYA J. T., VON HERZEN B. P.: Ray tracing volume densities. In *Proc. 11th annual conference on Computer graphics and interactive techniques* (1984), pp. 165–174. 2
- [Lev84] LEVERMORE C. D.: Relating Eddington factors to flux limiters. *Journal of Quantitative Spectroscopy and Radiative Transfer* 31 (1984), 149–160. 5
- [LK74] LARSEN E. W., KELLER J. B.: Asymptotic solution of neutron transport problems for small mean free paths. *Journal of Mathematical Physics* 15 (1974), 75–81. 1, 4, 5
- [LP81] LEVERMORE C. D., POMRANING G. C.: A flux-limited diffusion theory. *Astrophysical Journal* 248 (1981), 321–334. 1, 4, 5, 6
- [LSR*13] LI D., SUN X., REN Z., LIN S., TONG Y., GUO B., ZHOU K.: Transcut: Interactive rendering of translucent cutouts. *IEEE Transactions on Visualization and Computer Graphics* 19, 3 (2013), 484–494. 2
- [LW96] LAFORTUNE E. P., WILLEMS Y. D.: Rendering participating media with bidirectional path tracing. In *Proc. Conference on Rendering Techniques* (1996), Springer, pp. 91–100. 1
- [Min78] MINERBO G. N.: Maximum entropy Eddington factors. *Journal of Quantitative Spectroscopy and Radiative Transfer* 20 (1978), 541–545. 5
- [MSM*04] MAX N., SCHUSSMAN G., MIYAZAKI R., IWASAKI K., NISHITA T.: Diffusion and multiple anisotropic scattering for global illumination in clouds. *WSCG 12* (2004), 277–285. 2, 10
- [OAH00] OLSON G. L., AUER L. H., HALL M. L.: Diffusion, P_1 , and other approximate forms of radiation transport. *Journal of Quantitative Spectroscopy and Radiative Transfer* 64 (2000), 619–634. 5
- [Pom82] POMRANING G.: Flux limiters and Eddington factors. *Journal of Quantitative Spectroscopy and Radiative Transfer* 27, 5 (1982), 517–530. 5
- [QXF*07] QIU F., XU F., FAN Z., NEOPHYTOS N., KAUFMAN A., MUELLER K.: Lattice-based volumetric global illumination. *IEEE Transactions on Visualization and Computer Graphics* 13 (2007), 1576–1583. 2
- [RL86] RYBICKI G. B., LIGHTMAN A. P.: *Radiative Processes in Astrophysics*. Wiley, 1986. 2, 3
- [RSK06] RAAB M., SEIBERT D., KELLER A.: Unbiased global illumination with participating media. In *Proc. Monte Carlo and Quasi-Monte Carlo Methods*. Springer, 2006, pp. 591–605. 8
- [Sta95] STAM J.: Multiple scattering as a diffusion process. In *Proc. Conference on Rendering Techniques* (1995), pp. 41–50. 2, 3, 10
- [TS01] TURNER N. J., STONE J. M.: A module for radiation hydrodynamic calculations with ZEUS-2D using flux-limited diffusion. *The Astrophysical Journal Supplement Series* 135 (2001), 95–107. 1, 4
- [WZT*08] WANG J., ZHAO S., TONG X., LIN S., LIN Z., DONG Y., GUO B., SHUM H.-Y.: Modeling and rendering of heterogeneous translucent materials using the diffusion equation. *ACM Transactions on Graphics* 27, 1 (2008), 9. 2
- [YIC*10] YUE Y., IWASAKI K., CHEN B.-Y., DOBASHI Y., NISHITA T.: Unbiased, adaptive stochastic sampling for rendering inhomogeneous participating media. *ACM Transactions on Graphics* 29, 6 (2010), 177–185. 10
- [YZXW12] YAN L.-Q., ZHOU Y., XU K., WANG R.: Accurate translucent material rendering under spherical gaussian lights. *Computer Graphics Forum* 31, 7 (2012), 2267–2276. 2
- [ZM13] ZHANG Y., MA K.-L.: Fast global illumination for interactive volume visualization. In *Proc. Symposium on Interactive 3D Graphics and Games* (2013), pp. 55–62. 1